



Europäisches Patentamt
European Patent Office
Office européen des brevets



Publication number: **0 476 635 A2**

(12)

EUROPEAN PATENT APPLICATION

(21) Application number: **91115873.1**

(51) Int. Cl.⁵: **G06F 9/45**

(22) Date of filing: **18.09.91**

(30) Priority: **21.09.90 US 586778**

(43) Date of publication of application:
25.03.92 Bulletin 92/13

(84) Designated Contracting States:
DE FR GB IT

(71) Applicant: **Hewlett-Packard Company**
P.O. Box 10301 3000 Hanover Street
Palo Alto California 94303-0890(US)

(72) Inventor: **Lenkov, Dmitry**
1184 Hyde Avenue
San Jose, California(US)
Inventor: **Unni, Shankar**
2227 Peachtree Lane
San Jose, California(US)
Inventor: **Mehta, Michay**
1856 Deep Creek Ct.
San Jose, California(US)
Inventor: **McDowell, Mark W.**
1261 Solstice Lane
Ft. Collins, Colorado(US)
Inventor: **Dadoo, Manoj**
1650 Glenroy Drive
San Jose, California(US)
Inventor: **Melli, Bruno**
4030 Capstone Fort
Ft. Collins, Colorado(US)

(74) Representative: **Madgwick, Paul Roland et al**
Ladas & Parry Altheimer Eck 2
W-8000 München 2(DE)

(54) **Self identifying objects in computer systems.**

(57) The present invention is directed to computer software compilation systems and methods which support run-time data type identification of objects in computer programming languages which support polymorphism. The present invention comprises translators, compilers, and debuggers. The compiler and translator store modified debug information in an object code file. The modified debug information contains information regarding either multiple virtual tables or concatenated virtual tables. A debug lookup table is constructed from the modified debug information. The debugger uses the debug lookup table to determine the actual data types of the objects, and to completely and accurately display and modify the objects' contents. Also, innovative type inquiry operators reference the concatenated virtual tables to determine the actual data types of the objects during run-time. The operation of the compiler, translator, and debugger is transparent to computer programmers and operators. Therefore, the compiler, translator, and debugger support run-time data type identification of the objects in the computer programs in a user-friendly and error-free manner.

EP 0 476 635 A2

BACKGROUND OF THE INVENTION

The present invention relates to computer software compilation systems and methods which support run-time data type identification of instantiations of classes. In this patent document, instantiations of classes are also called objects.

Specifically, the present invention relates to computer software compilation systems and methods which support run-time data type identification of objects in computer programming languages which support polymorphism. Polymorphism is the ability to call a variety of functions using exactly the same interface. The present invention is adapted and intended to operate with computer programming languages (such as C++) which use derived classes and virtual functions to implement polymorphism, where virtual function calls are implemented as indirect function calls through tables of function addresses generated by compilers. Such tables are called virtual tables.

This patent document often refers to a C++ computer programming language to illustrate the features, structure, and operation of the present invention. This patent document also refers to the C++ computer programming language to describe the background of the present invention.

However, the references to the C++ computer programming language are included in this patent document for illustrative purposes only. The present invention is not limited to the C++ computer programming language. Rather, the present invention is adapted and intended to operate with computer programming languages which support polymorphism and which use virtual tables to store virtual member function addresses. The C++ computer programming language is an example of such a computer programming language.

Conventional C++ translators, compilers, and debuggers do not support run-time data type identification of objects in computer programs. In other words, accurate data type information for the objects in the computer programs which are produced using the conventional C++ translators and compilers is not available at run-time. Consequently, the conventional C++ debuggers are not adapted to extract and utilize the accurate, run-time data type information for the objects in the computer programs which are produced using the conventional C++ translators and compilers.

The unavailability of such data type information during run-time limits the functionality of the debuggers. For example, without access to such data type information during run-time, the debuggers cannot reliably refer to pointers and reference variables to determine and use the actual data types of the objects (where the pointers and reference variables point to the objects). Also, the debuggers cannot reliably use the pointers and reference variables to completely and accurately display the objects' contents.

Also, the availability of such data type information during run-time could be used as the basis for providing innovative type inquiry operators in languages such as C++. Such operators would be useful when writing code which needs to identify the exact data types of the objects.

The problem identified above is described below in greater detail with reference to Figures 1 and 2.

Referring first to Figure 1, suppose the statements in Table 1, below, are found in a C++ computer program.

TABLE 1

```

class ANIMAL {
public:
    virtual int x();
    virtual int y();
};

ANIMAL *dog_ptr = new ANIMAL;
ANIMAL *horse_ptr = new ANIMAL;

int i;

i = (*dog_ptr).x();

```

ANIMAL is declared as a class having two virtual member functions, x() and y(). A dog_ptr 136 and a

horse_ptr 134 are declared as pointers to type ANIMAL. The dog_ptr 136 and horse_ptr 134 are assigned to point to objects 124, 122, respectively, of type ANIMAL.

Figure 1 graphically presents an object code representation 102 of the class and pointer declarations and the object instantiations from Table 1. The object code representation 102 is generated by the
5 conventional C++ translators and compilers.

As shown in Figure 1, the object code representation 102 includes a member function x() code segment 106 and a member function y() code segment 108. The object code representation 102 also includes a virtual table 104 which is associated with the class ANIMAL. Generally, virtual tables are associated with classes that contain virtual member functions.

10 The virtual table 104 includes an entry 110, 112 for the member function y() and an entry 114, 116 for the member function x(). Generally, virtual tables associated with a class include one entry for each member function in the class.

The entry 110, 112 for the member function y() includes a field 112 that contains a pointer 118 to the member function y() code segment 108. The pointer 118 represents a virtual member function address of the member function y() code segment 108. The entry 110, 112 also includes a field 110 that contains
15 information related to calling the member function y() code segment 108.

The entry 114, 116 for the member function x() includes fields 114, 116 that are analogous to those of the entry 110, 112 for the member function y().

The objects 122, 124 of the type ANIMAL include self-descriptive information. The self-descriptive
20 information includes virtual table address fields 126, 128, respectively. The virtual table address fields 126, 128 contain virtual table addresses 130, 132, respectively, which point to the virtual table 104 associated with the class ANIMAL. Generally, objects of the same class contain virtual table addresses which point to the same virtual table.

In Table 1, the statement

25
i = (*dog_ptr).x();

calls the member function x() associated with the object 124 pointed to by the dog_ptr 136 (and also associated with the ANIMAL class) and assigns the integer returned by x() to a variable i. The member
30 function x() code segment 106 is accessed via the pointer 120 in the virtual table 104. The virtual table 104 is accessed via the virtual table address 132 in the object 124. The object 124 is accessed via the dog_ptr 136.

Referring now to Figure 2, suppose the statements in Table 2, below, are found in the C++ computer
35 program.

TABLE 2

```

40      class ONE {
          public:
              int I1;
              virtual int mem_1();
          };
45
      class TWO {
          public:
              int I2;
              virtual int mem_2();
50          };

      class THREE: public ONE, public TWO {
          int I3;
          };
55

      THREE *Ptr3 = new THREE;
```

ONE is declared as a class having a data member I1 and a virtual member function mem_1(). Similarly, TWO is declared as a class having a data member I2 and a virtual member function mem_2(). THREE is a derived class of the ONE and TWO classes and has a data member I3.

A Ptr3 226 is declared as a pointer to the class THREE and is assigned to point to an object 202 of the class THREE.

Figure 2 graphically presents an object code representation 102 of the statements from Table 2. Specifically, Figure 2 graphically presents an internal layout of the instantiated object 202 of the class THREE. The object code representation 102 is generated by the conventional C++ translators and compilers.

As shown in Figure 2, the object 202 includes the data members I1 206, I2 210, and I3 214. The object 202 also includes a virtual table address field 208 which contains a virtual table address 220 of a virtual table 216. The object 202 also includes a virtual table address field 212 which contains a virtual table address 222 of a virtual table 218. The virtual tables 216, 218 are similar to the virtual table 104 shown in Figure 1. Since the Ptr3 226 is declared as a pointer to the class THREE, the debugger and the innovative type inquiry operators can use the Ptr3 226 to determine the actual data type of the object 202. Also, since the Ptr3 226 points to the beginning of the object 202, the debugger can use the Ptr3 226 to completely and accurately display the contents of the object 202.

Instead of containing the statement,

THREE *Ptr3 = new THREE;

suppose Table 2 contained the statement,

ONE *Ptr1 = new THREE;

This new statement declares a Ptr1 224 as a pointer to the class ONE. However, this statement assigns the Ptr1 224 to point to the object 202 of the class THREE.

As shown in Figure 2, because Ptr1 224 points to the beginning of the object 202, the debugger can use the Ptr1 224 to completely and accurately display the contents of the object 202. However, because Ptr1 224 is declared as a pointer to the class ONE, the debugger and the innovative type inquiry operators cannot use the Ptr1 224 to determine the actual data type of the object 202.

Instead of containing the statement,

ONE *Ptr1 = new THREE;

suppose Table 2 contained the statement,

TWO *Ptr2 = new THREE;

This new statement declares a Ptr2 228 as a pointer to the class TWO. However, this statement assigns the Ptr2 228 to point to the object 202 of the class THREE.

As shown in Figure 2, because Ptr2 228 is declared as a pointer to the class TWO, the debugger and the innovative type inquiry operators cannot use the Ptr2 228 to determine the actual data type of the object 202. Also, because Ptr2 228 does not point to the beginning of the object 202, the debugger cannot use the Ptr2 228 to completely and accurately display the contents of the object 202.

Therefore, as illustrated in Figures 1 and 2 and as described above, the conventional C++ translators, compilers, and debuggers are inadequate in that they do not support run-time data type identification of objects in computer programs.

Prior approaches for solving such inadequacies of the conventional C++ translators, compilers, and debuggers involve explicit user intervention. Under these prior approaches, computer programmers are required to modify class declarations in the computer programs by adding additional members to the class declarations. The additional members are used to store the actual data types of the objects. Specifically, by operation of external tools, predefined virtual member functions, or the conventional C++ compiler, the actual data types of the objects are stored in the additional members. Consequently, data type identification of the objects during run-time is possible by referring to the additional members.

The prior approaches, however, are faulty in that they are not transparent to the computer programmers. Also, since the computer programmers are required to explicitly modify the class declarations, the prior approaches are laborious and error-prone.

SUMMARY OF THE INVENTION

The present invention is directed to software compilation systems and methods which support run-time data type identification of objects in computer programs.

5 Specifically, the present invention is directed to computer software compilation systems and methods which support run-time data type identification of objects in computer programming languages which support polymorphism. Polymorphism is the ability to call a variety of functions using exactly the same interface. The present invention is adapted and intended to operate with computer programming languages (such as C++) which use derived classes and virtual functions to implement polymorphism, where virtual
10 function calls are implemented as indirect function calls through tables of function addresses generated by compilers. Such tables are called virtual tables.

The present invention comprises improved translators, improved compilers, and improved debuggers.

The present invention comprises a first approach and a second approach for supporting run-time data type identification of the objects in the computer programs.

15 According to the first approach, the improved compiler and translator store modified debug information in an object code file. The modified debug information contains information regarding multiple virtual tables. A debug lookup table is constructed from the modified debug information. The improved debugger uses the debug lookup table to determine the actual data types of the objects, and to completely and accurately display the objects' contents. Also, the innovative type inquiry operators reference the debug lookup table
20 to determine the actual data types of the objects during run-time.

According to the second approach, which is an improvement upon the first approach in terms of efficiency, the improved compiler and translator create concatenated virtual tables. As with the first approach, the improved compiler and translator store modified debug information in the object code file. The modified debug information contains information regarding the concatenated virtual tables. As with the
25 first approach, the debug lookup table is constructed from the modified debug information. The improved debugger uses the debug lookup table to determine the actual data types of the objects, and to completely and accurately display the objects' contents. The innovative type inquiry operators reference the concatenated virtual tables to determine the actual data types of the objects during run-time.

The operation of the compiler, translator, and debugger according to the first and second approaches of
30 the present invention (as briefly described above) is transparent to computer programmers and operators. Therefore, the compiler, translator, and debugger support run-time data type identification of the objects in the computer programs in a user-friendly and error-free manner.

A better appreciation of these and other advantages and features of the present invention, as well as how the present invention realizes them, will be gained from the following detailed description and drawings
35 of the various embodiments, and from the claims.

BRIEF DESCRIPTION OF THE DRAWINGS

The present invention will be described with reference to the accompanying drawings, wherein:
40 Figure 1 shows a virtual table for base classes.
Figure 2 shows a virtual table for base classes and derived classes.
Figure 3 shows a structural diagram of the computer hardware and software environment in which the present invention operates.
Figure 4 shows a structural diagram/functional flowchart of a first embodiment of the present invention. In
45 Figure 4, rectangles represent software modules/processes and ovals represent the inputs and outputs of the software modules/processes.
Figure 5 shows a structural diagram/functional flowchart of a second embodiment of the present invention. In Figure 5, rectangles represent software modules/processes and ovals represent the inputs and outputs of the software modules/processes.
50 Figure 6 shows a structural diagram/functional flowchart of a compiler. In Figure 6, rectangles represent software modules/processes and ovals represent the inputs and outputs of the software modules/processes.
Figure 7 shows a structural diagram/functional flowchart of a debugger. In Figure 7, rectangles represent software modules/processes and ovals represent the inputs and outputs of the software
55 modules/processes.
Figure 8 shows a functional flowchart for generating modified debug information.
Figure 9 shows a high-level functional flowchart of the debugger.
Figure 10 shows a functional flowchart for accurately determining data types of objects during run-time.

Figure 11 shows a functional flowchart for accurately and completely determining the contents of objects during run-time.

Figure 12 shows multiple virtual tables and a corresponding debug lookup table.

Figure 13 shows a concatenated virtual table and a corresponding debug lookup table.

5 Figure 14 shows a functional flowchart for determining start addresses of concatenated virtual tables.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

1.0. Overview of the Present Invention

10

The present invention 328 is directed to software compilation systems and methods which support run-time data type identification of objects in computer programs 402.

Specifically, the present invention is directed to computer software compilation systems and methods which support run-time data type identification of objects in computer programming languages which support polymorphism. Polymorphism is the ability to call a variety of functions using exactly the same interface. The present invention is adapted and intended to operate with computer programming languages (such as C++) which use derived classes and virtual functions to implement polymorphism, where virtual function calls are implemented as indirect function calls through tables of function addresses generated by compilers. Such tables are called virtual tables.

20 The present invention 328 comprises improved translators 504, improved compilers 406, and improved debuggers 420.

As noted above, this patent document often refers to the C++ computer programming language to illustrate the features, structure, and operation of the present invention 328.

However, these references to the C++ computer programming language are included in this patent document for illustrative purposes only. The present invention 328 is not limited to the C++ computer programming language. Rather, the present invention 328 is adapted and intended to operate with computer programming languages which support polymorphism and which use virtual tables to store virtual member function addresses . The C++ computer programming language is an example of such a computer programming language.

30 As shown in Figure 3, the present invention is an application computer program 328 which operates on a computer platform 302. The computer platform 302 includes certain hardware units 308 including a central processing unit (CPU) 312, a random access memory (RAM) 310, and an input/output interface 314. The application computer program 328 may reside on the RAM 310. The computer platform 302 also includes an operating system 304, and may include microinstruction code 306. Various peripheral components may be connected to the computer platform 302, such as a terminal 318, a data storage device 322, and a printing device 326.

In a preferred embodiment of the present invention 328, the computer platform 302 is a HP9000 Series 300, 600, or 800 computer platform and the operating system 304 which runs thereon is HP-UX version 7.0. Also, the application computer program 328 of the present invention 328 is preferably written in the C++ computer programming language.

40 As shown in Figure 4, a first embodiment of the present invention 328 comprises the compiler 406, a linker 430, and the debugger 420.

As shown in Figure 5, a second embodiment of the present invention 328 comprises a translator/compiler combination 502, which includes the translator 504 and a compiler 506. The translator/compiler combination 502 performs the functions of the compiler 406 of the first embodiment. The structure and operation of the translator 504 and compiler 506 of the second embodiment is similar to that of the compiler 406 of the first embodiment. Therefore, the structure and operation of the translator 504 and compiler 506 of the second embodiment, with regard to the features and functions of the present invention 328 as taught and suggested in this patent document, should be obvious to persons with ordinary skill in the art based on the following detailed discussion of the structure and operation of the compiler 406 of the first embodiment.

Some aspects of the present invention 328 can be implemented using existing compiler, translator, linker, and debugger technology. However, modifications upon existing compiler, translator, and debugger technology are necessary to achieve the improvements of the present invention 328. The following discussions focus on these modifications upon existing compiler, translator, and debugger technology. For a general discussion of existing compiler, translator, linker, and debugger technology, see Compilers, Principles, Techniques, and Tools by Alfred V. Aho, Ravi Sethi, and Jeffrey D. Ullman (Addison Wesley 1986), which is herein incorporated in its entirety by reference.

The present invention 328 comprises a first approach and a second approach for supporting run-time data type identification of the objects in the computer programs 402.

According to the first approach, the improved compiler 406 stores modified debug information 416 in an object code file 412. The modified debug information 416 contains information regarding multiple virtual
5 tables 1208, 1216, 1220, 1224. A debug lookup table 1236 is constructed from the modified debug information 416. The improved debugger 420 uses the debug lookup table 1236 to determine the actual data types of the objects, and to completely and accurately display and modify the objects' contents.

Also according to the first approach, a lookup table that is similar to the debug lookup table 1236 can be generated. Innovative type inquiry operators can reference the lookup table to determine the actual data
10 types of the objects during run-time.

According to the second approach, which is an improvement upon the first approach, the improved compiler 406 creates concatenated virtual tables 1318. The improved compiler 406 stores the modified debug information 416 in the object code file 412. The modified debug information 416 contains information regarding the concatenated virtual tables 1318. The debug lookup table 1236 is constructed from the
15 modified debug information 416. The improved debugger 420 uses the debug lookup table 1236 to determine the actual data types of the objects, and to completely and accurately display and modify the objects' contents. The innovative type inquiry operators reference the concatenated virtual tables 1318 to determine the actual data types of the objects during run-time.

The operation of the compiler 406 and debugger 420 according to the first and second approaches of
20 the present invention 328 (as briefly described above) is transparent to computer programmers and operators. Therefore, the compiler 406 and debugger 420 support run-time data type identification of the objects in the computer programs 402 in a user-friendly and error-free manner.

As shown in Figure 4, and as noted briefly above, the first embodiment of the present invention 328 comprises the compiler 406, the linker 430, and the debugger 420. The compiler 406 receives the source
25 code 402 and generates object code 414 and the debug information 416. The object code 414 and the debug information 416 are stored in the object code file 412.

The linker 430 links the object code file 412 with other object code files 432 and forms an executable file 438. The executable file 438 includes at least the object code 414 and the debug information 416.

The debugger 420 receives the executable file 438 and, under the direction of a user (not shown in
30 Figure 4), uses the object code 414 and the debug information 416 in the executable file 438 to identify and correct errors in the source code 402.

A structural diagram/functional flowchart of the compiler 406 is shown in Figure 6. In Figure 6, rectangles represent software modules/processes and ovals represent the inputs and outputs of the software modules/processes.

As shown in Figure 6, the compiler 406 includes a scanner/parser 602, a semantic analyzer 610, a
35 debug information generator 620, and a code generator 624.

The scanner/parser 602 receives the source code 402 and builds an abstract syntax tree 606. The AST 606 essentially models the structure of the source code 402 being compiled. The AST 606 is composed of subtrees which represent declarations, function definitions, statements, and expressions.

The semantic analyzer 610 receives the AST 606 as input. The semantic analyzer 610 "decorates" the
40 nodes of the AST 606 with pointers to symbol table and type table entries (not shown in Figure 6) to produce a decorated abstract syntax tree 614.

The code generator 624 receives the decorated abstract syntax tree 614 as input. The code generator
1130 generates the object code 414 from the decorated abstract syntax tree 614 and places the object
45 code 414 in the object code file 412.

The debug information generator 620 also receives the decorated abstract syntax tree 614 as input. The
debug information generator 620 generates and stores modified debug information 416 in the object code
file 412. According to the first approach of the present invention, the modified debug information 416
50 contains information regarding the multiple virtual tables 1208, 1216, 1220, 1224. According to the second approach of the present invention, the modified debug information 416 contains information regarding the concatenated virtual tables 1318. The debug information generator 620 is described in detail below.

A structural diagram/functional flowchart of the debugger 420 is shown in Figure 7. In Figure 7, other
than the terminal 318 and a user 714, rectangles represent software modules/processes and ovals represent
the inputs and outputs of the software modules/processes.

As shown in Figure 7, the debugger 420 includes a preprocessor 702 and a main debugger 710.
55

The preprocessor 702 receives the executable file 438 (comprising the object code 414 and debug
information 416) and generates a preprocessed debug information 706. The preprocessed debug informa-
tion 706 is the same as the contents of the executable file 438, except the preprocessor 702 removes

duplicate information, fixes up pointers between various parts of the debug information 416, and collects certain data into several special tables which are added to the preprocessed debug information 706 for use by the main debugger 710. Specifically, the preprocessor 702 generates the debug lookup table 1236. The debug lookup table 1236 allows the main debugger 710 to accurately and completely determine at run-time the data type of objects referenced by pointers and reference variables.

The main debugger 710 receives the preprocessed debug information 706. Under the direction of a user 714, who communicates with the main debugger 710 via the terminal 318, the main debugger 710 uses the preprocessed debug information 706 to aid the user 714 in locating and correcting errors in the source code 402.

Specifically, the main debugger 710 uses the preprocessed debug information 706 to, among other things, display the source code 402, display an assembly code translation of the source code 402 generated by the compiler 406, to control and alter the flow of execution through the source code 402, to display variable types, and to view and set variable values.

To support data type identification of objects during run-time, the debugger 420 constructs the debug lookup table 1236. The debugger 420 uses the debug lookup table 1236 to determine the actual data types of the objects, to completely and accurately display the objects' contents, and to allow the user 714 to modify the objects' contents. The debugger 420 is discussed in detail below.

2.0. First Approach -- Using Multiple virtual Tables

This section describes the operation of the debug information generator 620 and the debugger 420 according to the first approach of the present invention 328. First, this section describes the format of the debug lookup table 1236.

Suppose the statements in Table 3 are contained in the source code 402. Then Figure 12 illustrates a corresponding debug lookup table 1236 according to the first approach of the present invention.

TABLE 3

```

class A {
public:
    // members
};

class B {
public:
    // members
};

class C: public A, public B {
    // members
};

A *ptr1 = new A;
B *ptr2 = new B;
B *ptr4 = new C;

```

A ptr1 1202 is declared as a pointer of class A and assigned to point to an object 1204 of class A. The object 1204 contains a virtual table address 1206 which points to a virtual table 1208.

A ptr2 1210 is declared as a pointer of class B and assigned to point to an object 1212 of class B. The object 1212 contains a virtual table address 1214 which points to a virtual table 1216.

A ptr4 1234 is declared as a pointer of class B and assigned to point to an object 1226/1228/1260 of class C. The object 1226/1228/1260 contains virtual table addresses 1218, 1222 which point to virtual tables 1220, 1224, respectively.

The debug lookup table 1236 contains an entry 1270, 1272, 1274, 1276 for every virtual table. The entries 1270, 1272, 1274, 1276 contain three fields: a class type identification field 1238, an object offset

field 1240, and a virtual table address field 1242.

The virtual table address field 1242 contains the virtual table addresses 1244, 1246, 1248, 1250 of the virtual tables 1208, 1216, 1224, 1220 associated with the entries 1270, 1272, 1274, 1276.

The class type identification field 1238 contains data type information of the objects associated with the virtual table addresses in the virtual table address fields 1242. For example, the debug lookup table 1236 entry 1270 stores a virtual table address 1244. The virtual table address 1244 points to a virtual table 1208. The virtual table 1208 is associated with the class A. Therefore, the class type identification field 1238 associated with the entry 1270 identifies the class A.

The object offset field 1240 contains pointer object offsets to the beginning of objects. As shown in Figure 12, for example, the ptr4 1234 does not point to the beginning of the object 1226/1228/1260. Specifically, the ptr4 1234 is displaced by an object offset of d (d is equal, in this case, to the size of an object of class A) from the beginning of the object 1226/1228/1260. Therefore, the object offset field 1240 of entry 1276 (the entry 1276 stores a virtual table address 1250 which is equal to a virtual table address 1218 associated with the ptr4 1234) contains a value equal to the offset d.

The operation of the debug information generator 620 and the main debugger 710 according to the first approach of the present invention 328 is described below with reference to Figures 8-12.

Figure 8 illustrates the manner in which the debug information generator 620 generates the modified debug information 416 according to the first approach of the present invention 328.

In Step 806, the debug information generator 620 locates the classes to which the objects 1204, 1212, 1226/1228/1260 belong in the decorated AST 614. The debug information generator 620 then determines locations 1280, 1282, 1286, 1284 in the objects 1204, 1212, 1226/1228/1260 where the virtual table addresses 1206, 1214, 1218, 1222 are stored.

In Step 810, the debug information generator 620 stores the locations 1280, 1282, 1286, 1284 in the modified debug information 416 (these locations are offsets inside the objects where pointers to the virtual table are stored).

In Step 814, the debug information generator 620 reads the virtual table addresses 1206, 1214, 1218, 1222 from the objects 1204, 1212, 1226/1228/1260. The debug information generator 620 also determines the object offsets. As noted above (with reference to Figure 12), the object offsets represent pointer offsets to the beginning of objects.

In Step 818, the debug information generator 620 stores the virtual table addresses 1206, 1214, 1218, 1222 and the object offsets in the modified debug information 416.

As indicated by Steps 822 and 828, Steps 806, 810, 814, and 818 are repeated for all objects which are instantiations of classes containing virtual functions.

Figure 9 presents a functional flowchart of the high-level operation of the debugger 420 according to the first approach of the present invention.

In Step 906, the debugger 420 uses the modified debug information 416 to create the debug lookup table 1236.

In Step 910, the debugger 420 receives user requests. The user may request the debugger 420 to display the source code 402, display the assembly code translation of the source code 402 generated by the compiler 406, to control and alter the flow of execution through the source code 402, to display variable types, and to view and set variable values.

In Step 914, the debugger 420 processes the user requests. In performing Step 914, the debugger 420 refers to the debug lookup table 1236 when it must determine the data types of objects, accurately and completely display the contents of objects, or allow the user 714 to modify the contents of the objects.

The manner in which the debugger 420 accurately determines the data types of objects is presented in Figure 10. The flowchart in Figure 10 is described below with reference to the ptr4 1234 in Figure 12. Specifically, suppose the user requests that the debugger 420 display the data type of the object 1226/1228/1260 pointed to by the ptr4 1234.

In Step 1006, the debugger 420 reads the virtual table address 1286 in the object 1226/1228/1260 referenced by the ptr4 1234.

In Step 1010, the debugger 420 searches the debug lookup table 1236 for an entry whose virtual table address (in the virtual table address field 1242) is equal to that read in Step 1006. For the present example, the debugger 420 locates the entry 1276. Then, the debugger 420 reads the class type identification field 1238 of the entry 1276 to determine the data type of the object 1226/1228/1260 pointed to by the ptr4 1234. For the present example, the debugger 420 correctly determines that the data type is class C.

In Step 1014, the debugger 420 uses the data type information to finish the processing of the user request. For the present example, the debugger 420 may simply display the data type information.

The flowchart in Figure 10, for determining the data types of objects referenced by pointers or

reference variables, is also applicable for the innovative type inquiry operators. For example, the compiler 406 could emit a lookup table 1236 to the object code 414. Then, the innovative type inquiry operators could refer to the lookup table 1236 (according to the flowchart in Figure 10) to determine the data type of objects during run-time.

6 The manner in which the debugger 420 accurately and completely determines the contents of objects is presented in Figure 11. The flowchart in Figure 11 is described below with reference to the ptr4 1234 in Figure 12. Specifically, suppose the user requests that the debugger 420 display the contents of the object 1226/1228/1260 pointed to by the ptr4 1234.

10 In Step 1106, the debugger 420 reads the virtual table address 1286 in the object 1226/1228/1260 that is referenced by the ptr4 1234.

In Step 1110, the debugger 420 searches the debug lookup table 1236 for an entry whose virtual table address (in the virtual table address field 1242) is equal to that read in Step 1106. For the present example, the debugger 420 locates the entry 1276. Then, the debugger 420 reads the object offset field 1240 of the entry 1276. For the present example, the debugger 420 reads the object offset d. As shown in Figure 12, 15 the offset d represents the object offset of the ptr4 1234 from the beginning of the object 1226/1228/1260.

In Step 1114, the debugger 420 calculates a start address of the object 1226/1228/1260. In the preferred embodiment, the debugger 420 calculates the start address by subtracting the object offset d from the ptr4 1234.

20 In Step 1118, the debugger 420 uses the start address to display the contents of the object 1226/1228/1260 to the user.

If, instead, the user 714 had requested that the contents of the object 1226/1228/1260 be modified, all steps in Figure 11 would be identical to the user request for displaying the object contents except for Step 1118, which would be replaced by a Step 1118' (not explicitly shown in Figure 11). In Step 1118', the debugger 420 uses the start address to modify the contents of the object 1226/1228/1260.

25 3.0. Second Approach – Using Concatenated Virtual Tables

This section describes the operation of the debug information generator 620 and the main debugger 710 according to the second approach of the present invention 328. As noted above, the second approach 30 is an improvement of the first approach.

As noted above, the first approach enables the debugger 420 to accurately determine the data type of objects, to accurately and completely display the contents of the objects, and to allow the user 714 to modify the contents of the objects. Also, the first approach enables the innovative type inquiry operators to accurately determine the data type of objects during run-time.

35 As a general rule, however, table lookups during run-time should be avoided for efficiency reasons. Therefore, while achieving the objectives of the present invention 328, the first approach is relatively inefficient with regard to enabling the innovative type inquiry operators to determine the data type of objects during run-time.

40 The second approach of the present invention 328 solves the efficiency problem of the first approach by using the concatenated virtual tables 1318, rather than the multiple virtual tables 1208, 1216, 1220, 1224. Specifically, a unique, distinct concatenated virtual table 1318 is generated for each class. For some classes, the concatenated virtual table 1318 may be made up from only one virtual table, but for other classes, the concatenated virtual table 1318 may contain multiple virtual tables.

45 This section describes the second approach of the present invention 328 by first describing the format of the concatenated virtual tables 1318.

Suppose the statements in Table 4 are contained in the source code 402. Figure 13 illustrates a corresponding concatenated virtual table 1318 and a debug lookup table 1236 according to the second approach of the present invention 328.

50

55

TABLE 4

```

5      class A {
      public:
          // members
      };

10     class B {
      public:
          // members
      };

15     class C {
      public:
          // members
      };

20     class D: public A, public B, public C {
          // members
      };

25     B *ptr2 = new D;

```

A ptr2 1336 is declared as a pointer of class B and assigned to point to an object 1390 of class D. The object 1390 contains virtual table addresses 1312, 1314, 1316 that point to virtual tables 1360, 1362, 1364, respectively.

30 Generally, compilers 406 emit multiple virtual tables for an object. This is shown in Figure 12, where the compiler 406 emitted the multiple virtual tables 1220, 1224 for the object 1226/1228/1260. In this patent document, the term "multiple" means that the virtual tables are not contiguous in memory.

Generally, the multiple virtual tables 1220, 1224 essentially comprise two fields, a field 112, 116 containing a pointer to a member function code segment and a field 110, 114 containing information related to calling the member function code segment. This is shown in Figure 1.

35 According to the second approach of the present invention 328, the compiler 406 emits a concatenated virtual table 1318 for an object 1390 of class D. This is shown in Figure 13, where the compiler 406 emitted the concatenated virtual table 1318 for the object 1390. The concatenated virtual table 1318 contains virtual tables 1360, 1362, 1364. In this patent document, the term "concatenated" means that the virtual tables are contiguous in memory.

Also according to the second approach of the present invention 328, the virtual tables 1360, 1362, 1364 emitted by the compiler 406 contain three fields 1306, 1308, 1310. The field 1306 corresponds to the field 112, 116 containing the pointer to the member function code segment. The field 1310 corresponds to the field 110, 114 containing information related to calling the member function code segment.

45 The additional field 1308 contains pointer table offset values. A pointer table offset value associated with a particular virtual table represents a displacement from the beginning of the virtual table to the beginning of a concatenated virtual table for which the virtual table is a part. For example, the pointer table offset value for the virtual table 1364 is 12, because the beginning of the virtual table 1364, pointed to by virtual table address 1316, is displaced 12 from the beginning of the concatenated virtual table 1318, pointed to by the virtual table address 1302. Note that the pointer table offset values in the table offset fields 1308 are distinct from the pointer object offset values in the object offset fields 1240 in the debug lookup table 1236.

The operation of the compiler 406 and the main debugger 710 according to the second approach of the present invention 328 is described below.

55 As noted above, the compiler 406 emits concatenated virtual tables 1318, rather than multiple virtual tables 1220, 1224, for objects according to the second approach of the present invention 328. The virtual tables 1360, 1362, 1364 contained in the concatenated virtual tables 1318, and emitted by the compiler 406, contain the table offset field 1308 in addition to the fields 1306 and 1310.

Also, the debug information generator 620 generates the modified debug information 416. The manner

in which the debug information generator 620 generates the modified debug information 416 is the same for both the first and second approaches of the present invention 328. Therefore, the operation of the debug information generator 620 according to the second approach of the present invention 328 is adequately described above with reference to Figure 8.

5 The manner in which the innovative type inquiry operators accurately determine the data type of objects during run-time is described below with reference to Figures 13 and 14. The flowchart in Figure 14 is described below with reference to the ptr2 1336 in Figure 13. Specifically, suppose a C++ operator, as part of its function, must determine the data type of the object 1390 pointed to by the ptr2 1336.

10 An innovative "typeof()" operator is an example of such a C++ type inquiry operator. The typeof() operator receives a pointer (or a reference variable) to an object and returns a data type indication of the object's data type. The typeof() operator returns the same data type indication for all pointers associated with the object. In Figure 13, for example, the typeof() operator returns the virtual table address 1302 for the ptr1 1332, ptr2 1336, ptr3 1338, and ptr4 1334. In this particular example, the data type indication is the virtual table address 1302 of the start of the concatenated virtual table 1318.

15 Referring now to Figure 14, in Step 1406, the C++ operator or function, such as typeof(), reads the virtual table address 1314 in the object 1390 referenced by the ptr2 1336.

In Step 1410, the C++ operator or function, such as typeof(), uses the virtual table address 1314 to access the virtual table 1362 and read the table offset field 1308. In the present example, the table offset value l1 is read.

20 In Step 1414, the C++ operator or function, such as typeof(), calculates the start address of the concatenated virtual table 1318. The start address of the concatenated virtual table 1318 is equal to the virtual table address 1302. In the preferred embodiment of the present invention 328, the start address is calculated by subtracting the table offset value l1 from the virtual table address 1314.

In Step 1418, the C++ operator or function retains the start address that was calculated in Step 1414 for further processing. For example, the typeof() operator may simply return the start address of the concatenated virtual table 1318.

The operation of the debugger 420 is the same for both the first and second approaches of the present invention 328. Therefore, the operation of the debugger 420 according to the second approach of the present invention 328 is adequately described above with reference to Figures 9, 10, and 11.

30 Note that the first and second approaches of the present invention 328 are applicable only to classes which contain virtual member functions. However, it should be obvious to those with ordinary skill in the art that compile-time switches could be implemented to emit pointers to unique tables inside all structures (except for those structures which are explicitly declared with 'extern "C"'). Such compile-time switches would ensure that the present invention 328 would apply to all classes.

35 Based on the discussion above, it should be obvious to those with ordinary skill in the art that the concatenated virtual tables 1318 are useful for implementing other innovative operators and functions. For example, by an obvious extension of the concatenated virtual tables 1318, a "sizeof()" operator for objects referenced by pointers and reference variables could be implemented. The sizeof() operator would receive pointers (or reference variables) to objects and return the actual size of the object, as opposed to the size of the class specified in the declaration of the pointer or reference variable.

Claims

1. A computer-based software compilation system, adapted for use with a source code, wherein the source code is written in a computer programming language which uses virtual tables to store virtual member function addresses, and wherein the source code includes objects having object data types, object contents, and virtual table addresses, the system comprising:

45 (1) first means for generating an object code and a debug information of the source code, said debug information comprising said virtual table addresses and virtual table address locations; and
50 (2) second means, adapted to operate under the direction of a user, for locating and correcting errors in the source code by using said debug information, said second means comprising:

(a) means for accurately determining said object data types; and

(b) means for accurately and completely determining said object contents;

55 wherein said first and second means are adapted to support object data type identification during run-time.

2. The system of claim 1, wherein said first means comprises:

(1) a scanner/parser for receiving the source code and generating an abstract syntax tree;

- (2) a semantic analyzer for receiving said abstract syntax tree and generating a decorated abstract syntax tree;
- (3) a code generator for receiving said decorated abstract syntax tree and generating said object code; and
- 5 (4) debug information generator means for receiving said decorated abstract syntax tree and generating said debug information.
3. The system of claims 1 or 2, wherein said second means comprises means for generating a debug lookup table, wherein said debug lookup table comprises entries for said virtual tables, and wherein
- 10 said entries comprise object data type identification fields, object offset fields, and virtual table address fields.
4. The system of claim 3, wherein said means for accurately determining said object data types comprises:
- 15 (1) means for reading said virtual table addresses by using pointers and reference variables, wherein said pointers and reference variables refer to said objects;
- (2) means for accessing said debug lookup table and reading data type identification information in said data type identification fields according to said virtual table addresses, wherein said data type identification information identifies said object data types; and
- 20 (3) means for using said data type identification information to process user requests.
5. A computer-based software compiler, adapted for use with a source code, wherein the source code is written in a computer programming language which uses virtual tables to store virtual member function addresses, and wherein the source code includes objects having object data types, object contents, and virtual table addresses, the compiler comprising:
- 25 (1) first means for receiving the source code and generating an abstract syntax tree;
- (2) second means for receiving said abstract syntax tree and generating a decorated abstract syntax tree;
- (3) third means for receiving said decorated abstract syntax tree and generating an object code of the source code; and
- 30 (4) debug information generator means for receiving said decorated abstract syntax tree and generating a debug information of the source code, wherein said debug information comprises said virtual table addresses and virtual table address locations;
- wherein said compiler is adapted to support object data type identification during run-time.
- 35 6. The system of claims 2 or 5, wherein said (fourth) debug information generator means comprises:
- (1) means for determining said virtual table address locations;
- (2) means for determining said virtual table addresses and object offsets; and
- (3) means for storing said virtual table address locations, virtual table addresses, and object offsets
- 40 in said debug information.
7. The system of claim 5, wherein said compiler further comprises:
- (1) means for generating a lookup table, wherein said lookup table comprises entries for said virtual tables, and wherein said entries comprise object data type identification fields, object offset fields, and virtual table address fields; and
- 45 (2) means for storing said lookup table in said object code;
- wherein said lookup table is adapted to be accessed by type inquiry operators, such that said type inquiry operators can determine said object data types during run-time.
- 50 8. The system of claims 5 or 7, wherein said compiler further comprises;
- (1) means for generating concatenated virtual tables, wherein said concatenated virtual tables comprise table offset fields; and
- (2) means for storing said concatenated virtual tables in said object code.
- 55 9. The system of claim 8, further adapted to support a typeof() operator for determining said object data types during run-time, said typeof() operator comprising:
- (1) means for reading said virtual table addresses by using pointers and reference variables, wherein said pointers and reference variables refer to said objects;

(2) means for accessing said concatenated virtual tables and reading pointer table offset values in said table offset fields according to said virtual table addresses;

(3) means for using said pointer table offset values to calculate concatenated virtual table starting addresses, wherein said concatenated virtual table starting addresses identify said data types; and

(4) means for returning said concatenated virtual table starting addresses.

10. A computer-based debugger, adapted to operate under the direction of a user, for locating and correcting errors in a source code by using a debug information, wherein the source code is written in a computer programming language which uses virtual tables to store virtual member function addresses, and wherein the source code includes objects having object data types, object contents, and virtual table addresses, and wherein said debug information comprises said virtual table addresses and virtual table address locations, the debugger comprising:

(1) first means for generating a debug lookup table, said debug lookup table comprising entries for said virtual tables, wherein said entries comprise object data type identification fields, object offset fields, and virtual table address fields;

(2) second means for accurately determining said object data types; and

(3) third means for accurately and completely determining said object contents;

wherein said debugger is adapted to support object data type identification during run-time.

11. The system of claim 10, wherein said second means comprises:

(1) means for reading said virtual table addresses by using pointers and reference variables, wherein said pointers and reference variables refer to said objects;

(2) means for accessing said debug lookup table and reading data type identification information in said data type identification fields according to said virtual table addresses, wherein said data type identification information identifies said object data types; and

(3) means for using said data type identification information to process user requests.

12. The system of claims 10 or 11, wherein said third means comprises:

(1) means for reading said virtual table addresses by using pointers and reference variables, wherein said pointers and reference variables refer to said objects;

(2) means for accessing said debug lookup table and reading pointer object offset values in said object offset fields according to said virtual table addresses;

(3) means for using said pointer object offset values to calculate object starting addresses; and

(4) means for using said object starting addresses to display and modify said object contents.




FIG. 1




FIG. 2

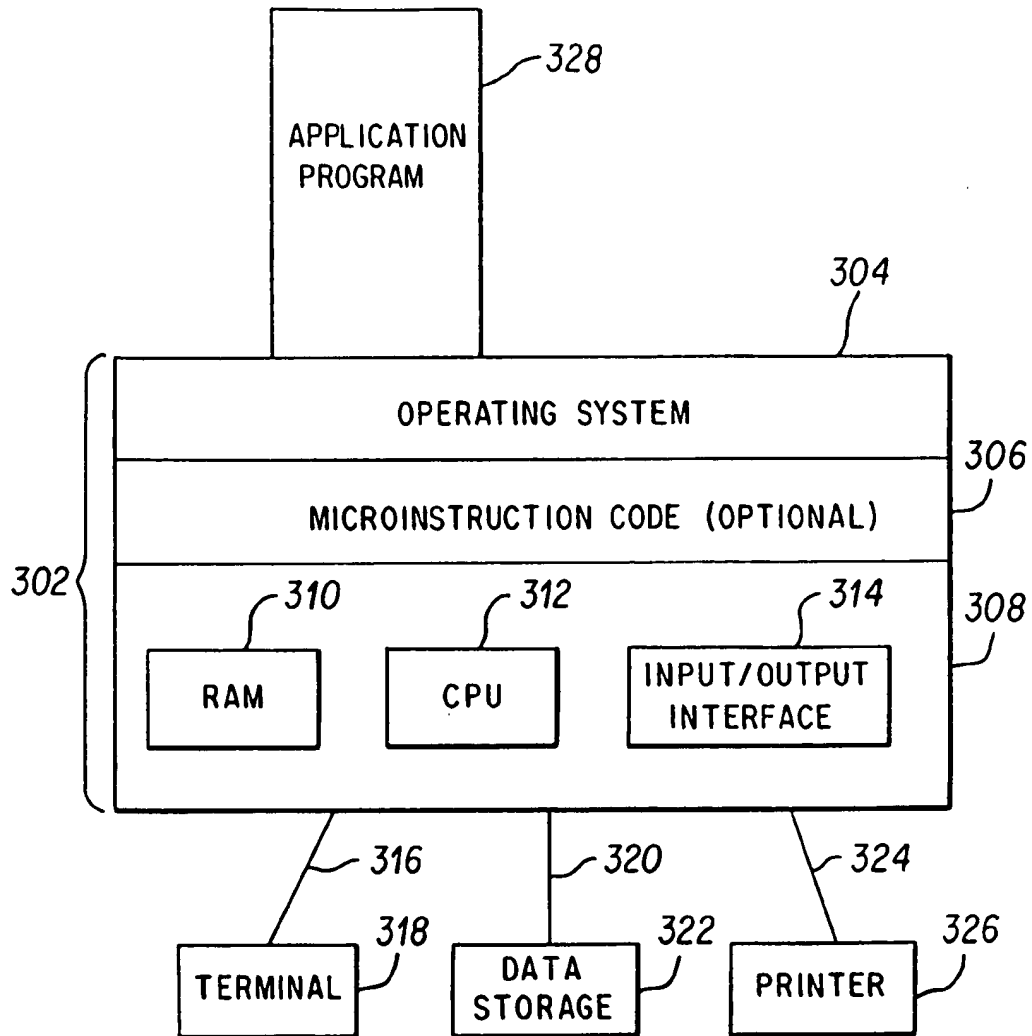


FIG. 3




FIG. 4




FIG. 5

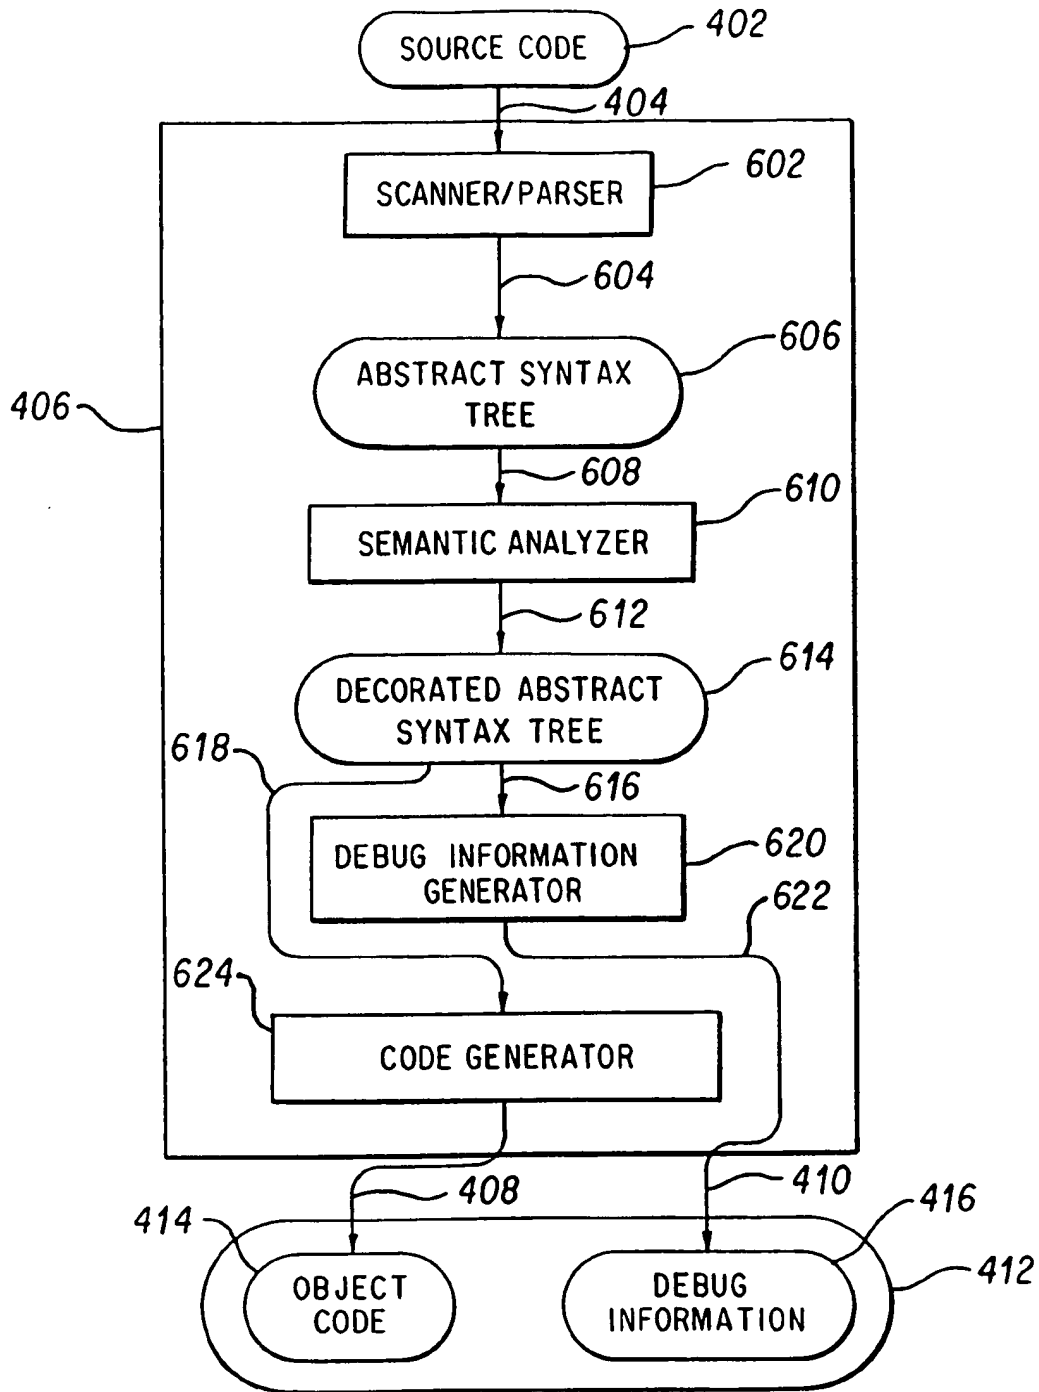


FIG. 6




FIG. 7

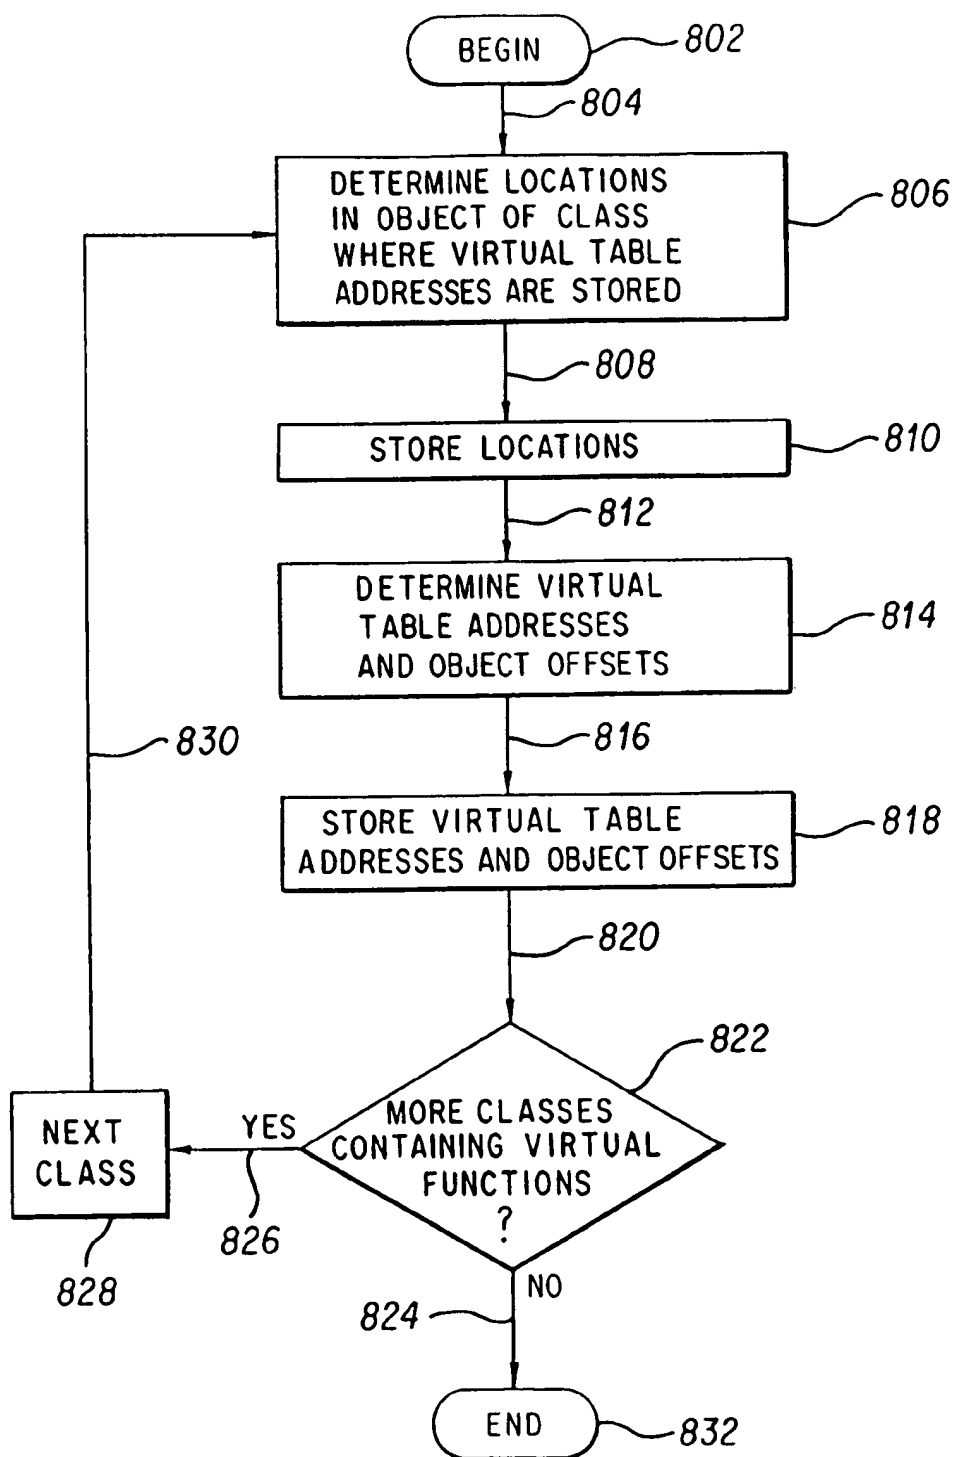


FIG. 8

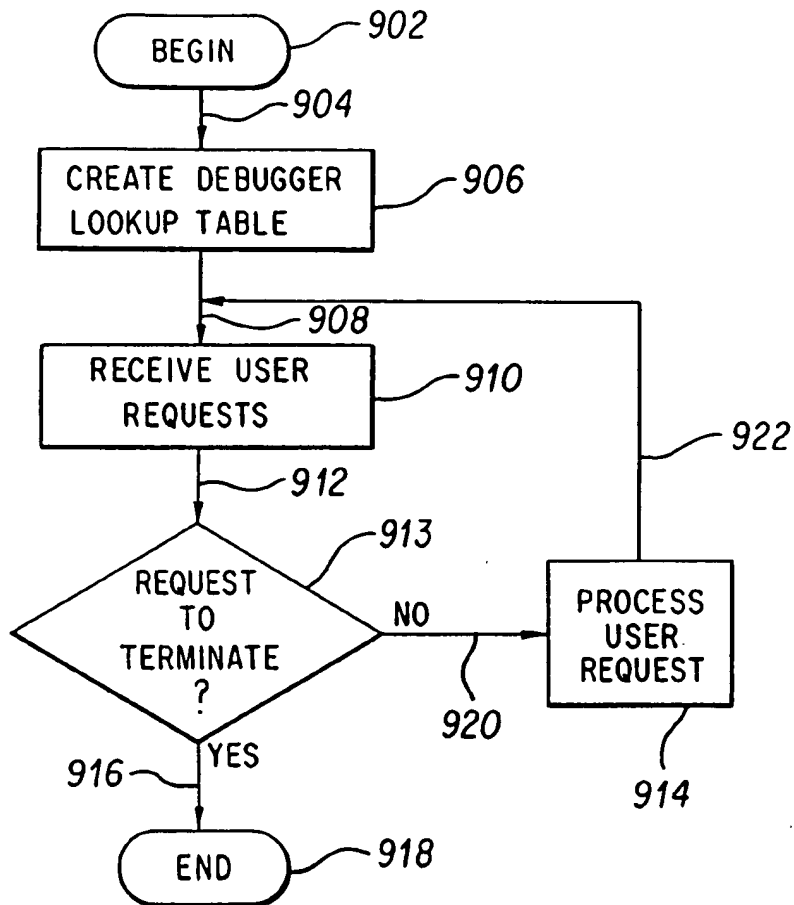


FIG. 9

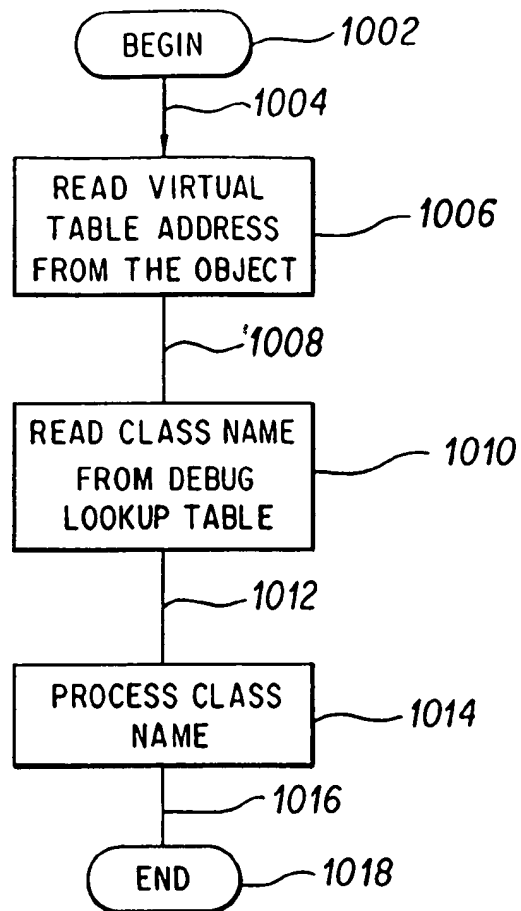


FIG. 10

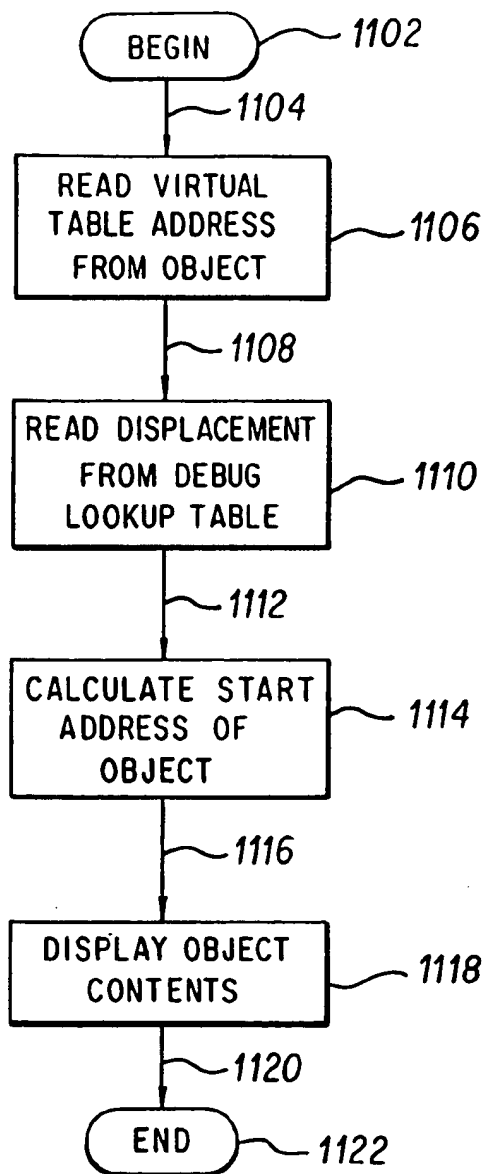


FIG. 11





FIG. 12






FIG. 14